

CLAIMS

We claim:

1. In a database system, a method of maintaining a self-tuning histogram having a plurality of existing buckets arranged in a hierarchical manner and defined by at least two bucket boundaries, a bucket volume, and a bucket frequency comprising the step of creating at least one new bucket in response to a query on the database wherein each new bucket is contained within at least one existing bucket and wherein the new bucket becomes a child bucket and the existing bucket becomes a parent bucket.
2. The method of claim 1 wherein each bucket has a rectangular shape.
3. The method of claim 1 wherein the boundaries of each new bucket correspond to a region of the database accessed by the query and the frequency of each new bucket is a number of data records returned by the query.
4. The method of claim 1 wherein a total number of buckets is limited to a predetermined budget.
5. The method of claim 4 comprising the step of merging buckets based on a merge criterion when the total number of buckets exceeds the predetermined budget.
6. The method of claim 5 wherein the merge criterion is a similar bucket density, wherein bucket density is based on the bucket frequency divided by the bucket volume.
7. The method of claim 1 further comprising the step of shrinking the boundaries of each new bucket if the boundaries of the new bucket intersect any existing bucket boundaries.

8. The method of claim 1 wherein the frequency of the parent bucket is diminished by the frequency of the child bucket.

9. In a database system, a method of maintaining a self-tuning histogram having a plurality of existing buckets arranged in a hierarchical manner and defined by at least two bucket boundaries, a bucket volume, and a bucket frequency comprising the steps of:

- a) examining the results of a query executed on the database;
- b) creating at least one candidate hole in the histogram based on the results of the query;
- c) modifying each candidate hole such that the modified hole is completely contained within at least one existing parent bucket and does not partially intersect any existing bucket; and
- d) creating a new child bucket in the histogram corresponding to each modified hole.

10. The method of claim 9 wherein each bucket has a rectangular shape.

11. The method of claim 9 wherein the boundaries of each candidate hole correspond to a region of the database accessed by the query and the frequency of each candidate hole is a number of data records returned by the query.

12. The method of claim 9 wherein a total number of buckets is limited to a predetermined budget.

13. The method of claim 12 comprising the step of merging buckets based on a merge criterion when the total number of buckets exceeds the predetermined budget.

14. The method of claim 13 wherein the merge criterion is a similar bucket density, wherein bucket density is based on the bucket frequency divided by the bucket volume.

15. The method of claim 9 wherein the frequency of the parent bucket is diminished by the frequency of the child bucket.

16. A computer readable medium having computer executable instructions for performing steps for maintaining a self-tuning histogram having a plurality of existing buckets arranged in a hierarchical manner and defined by at least two bucket boundaries, a bucket volume, and a bucket frequency, the steps comprising:

- a) examining the results of a query executed on the database;
- b) creating at least one candidate hole in the histogram based on the results of the query;
- c) modifying each candidate hole such that the modified hole is completely contained within at least one existing parent bucket and does not partially intersect any existing bucket; and
- d) creating a new child bucket in the histogram corresponding to each modified hole.

17. The computer readable medium of claim 16 wherein each bucket has a rectangular shape.

18. The computer readable medium of claim 16 wherein the boundaries of each candidate hole correspond to a region of the database accessed by the query and the frequency of each candidate hole is a number of data records returned by the query.

19. The computer readable medium of claim 16 comprising the step of merging buckets having a similar bucket density when the total number of buckets exceeds a predetermined budget.

20. An apparatus for maintaining a self-tuning histogram having a plurality of existing buckets arranged in a hierarchical manner and defined by at least two bucket boundaries, a bucket volume, and a bucket frequency comprising:

- a) means for examining the results of a query executed on the database;
- b) means for creating at least one candidate hole in the histogram based on the results of the query;
- c) means for modifying each candidate hole such that the modified hole is completely contained within at least one existing parent bucket and does not partially intersect any existing bucket; and
- d) means for creating a new child bucket in the histogram corresponding to each modified hole.

21. The apparatus of claim 20 wherein the means for creating at least one candidate hole assigns the boundaries of the candidate hole to correspond to a region of the database accessed by the query and the frequency of each candidate hole is assigned a number of data records returned by the query.

22. An apparatus for maintaining a self-tuning histogram having a plurality of existing buckets arranged in a hierarchical manner and defined by at least two bucket boundaries, a bucket volume, and a bucket frequency comprising:

- a) a memory device for storing a database comprising multiple data records;
- b) a computer having one or more processing units for executing a stored computer program, said computer including a rapid access memory store; and
- c) an interface for coupling the memory device for storing the database to the computer to allow records to be retrieved from the database; wherein
- d) the stored program has components including i) a component for examining the results of a query executed on the database; ii) a component for creating at least one candidate hole in the histogram based on the results of the query; iii) a component for modifying each candidate hole such that the modified hole is completely contained within at least one existing parent bucket and does not partially intersect any existing bucket; and iv) a component for creating a new child bucket in the histogram corresponding to each modified hole.

23. The apparatus of claim 22 wherein the component for creating at least one candidate hole assigns the boundaries of the candidate hole to correspond to a region of the database accessed by the query and the frequency of each candidate hole is assigned a number of data records returned by the query.